

The Marketer's Dilemma

or

The Tragedy of the Twitter Commons

Being an excerpt from Chapter 10 of *Social Media Marketing: A Game Theory Perspective*, to be published by Springer in June 2010.

• • •

Many years ago, *Saturday Night Live* aired a talk show spoof called "Ruined for Everyone," in which guests described how they were the first person to violate certain customs that were based on common trust, thereby

In the Volunteer's Dilemma, the participant must decide whether to make a short-term sacrifice in order to preserve a collective good that they themselves participate in.

ruining those customs for the rest of us. One of the guests was the first person to pump their gas and drive off without paying. Another guest was the first to trash a public restroom, and another ruined hitchhiking. The sketch was funny because we really do rue the loss of these simple trust-based systems, and we see them being threatened every day. The person who cuts ahead in the bakery line threatens the whole system, because the people who wait patiently in line may decide that they are suckers for sticking to the rules and defect, causing a run on the pastry counter. The problem is exacerbated in large systems like corporations, where collective rather than individual responsibility makes rule-breaking easier; a competitor who breaks the marketing rules with pernicious tactics and gets short-term results

creates a temptation for others to do the same.

This problem of how to get participants to behave themselves in a free and open system with no central governing authority is one that I have taken up in previous chapters in describing the free rider scenario. This same problem occurs in collective scenarios, in which an individual can gain more in the short term by defecting from the group's collective interest, even though doing so ruins the group's long-term interests and hurts the individual as well. In game theory this problem has been called the Volunteer's Dilemma, which, like its close cousin the Iterated Prisoner's Dilemma, is mainly concerned with how to enforce cooperation wherever it serves participants' mutual long-term interests.

In the Volunteer's Dilemma, the participant must decide whether to make a short-term sacrifice in order to preserve a collective good that they themselves participate in. That sacrifice will cost the individual more in the short-term, and therein lies the dilemma: it is tempting to act as a free rider for short-term gain, especially since making a sacrifice is no guarantee that others will make the same sacrifice. Minding your place in line in the bakery is an apt example, but game theory offers more dramatic ones. The most famous is the "Tragedy of the Commons," first articulated in an article by Garret Hardin in *Science* in 1968. Hardin invoked a 19th century philosophical tract contemplating the problem of overgrazing on lands held in common. When a herdsman decides whether to add another animal to his grazing flock, he reasons that the positive consequences – increased revenue – will accrue to him alone, while the negative consequences – overgrazing – will be shared in common with the other herdsman. He therefore rationally decides to add more animals, and could continue to do so into infinity, even though the resource itself is tragically finite. His fellow herdsman could be expected to follow the same logic and the same tragic course, until the common grazing ground is destroyed. As Hardin eloquently describes this outcome, "Ruin is the destination toward which all men rush, each pursuing his own best interest in a society that believes in the freedom of the commons. Freedom in a commons brings ruin to all" (Hardin 1968).

In contemplating marketers' use of free social network resources like Facebook, YouTube, and Twitter, one might object that these resources are altogether different from shared grazing lands, in that they are as infinite as their owners' willingness to continue adding server capacity. In theory, bad behavior by some marketers will not remove or even reduce the ability of other marketers to get their message out in these forums, since access is, to date, unlimited. But as you have probably surmised, server capacity is not the finite resource at risk of overgrazing; the finite resource is the consumer's attention.

Consumer attention is
the finite resource at
risk of overgrazing.

Social Media's Capacity Problem

Before considering how marketers will deal with the problem of the diminishing supply of consumer attention, I feel compelled to offer evidence of the problem itself:

- In June 2009, Nielsen reported that Twitter usage had grown 2,000% in one year's time. This anniversary coincided with what some predicted would be a meltdown of all Twitter-related applications, as the number of posts surpassed 2.1 billion, the maximum number of entries that most databases are equipped to handle (based on a 32-bit signed integer). Twitter has since been recoded to handle the additional capacity, and as of this writing, nearly 10 billion tweets have been posted.

- While the growth in the number of new blogs being created has slowed (presumably because everyone has one now), Technorati reports that there are approximately 900,000 new blog posts added to the blogosphere every 24 hours.
- Of course, users may reduce their Web usage in order to spend more time on their mobile devices, but they won't escape content saturation there; the number of mobile applications is expected to reach 100,000 by the end of 2010, and 10 million by 2020. The number of text messages sent in the U.S. alone in 2008 is estimated at 95 billion.
- With so much social content available, perhaps users can rely on expert guidance – a social media guru – to help sort it all out. But one must first sort out the gurus: there are 5,855 self-proclaimed social media experts on Twitter alone (Ochman 2009).

I share these numbers to give some sense of the scale of competition for consumers' attention; that competition, in turn, creates the overgrazing problem. It would be possible, in fact, to plot out this overgrazing mathematically, given the right data set. For instance, Nielsen's most recent report of Web usage shows that the average Facebook user spends a bit over 7 hours per month on the network. Suppose that average user is connected to 5 brands through Facebook, and they spend 1 out of 7 of their Facebook hours reading updates from those brands.

As the user increases the number of brands that they connect with, the average attention given to each brand declines.

The problem for the brand is that their allotted attention from that user is likely to decrease as the user's Facebook usage increases. As the user increases the number of brands that they connect with, the average attention given to each brand declines. The user may increase their total time on Facebook, but they cannot increase it indefinitely, and time spent on brands will inevitably lose out to time spent on friends and family as the user increases their number of personal connections. The brand continues to push updates, offering links to discounts and other costly enticements, but the user's attention still declines, as it must, which in turn makes the brand more aggressive. We are in the familiar death spiral of mutual defection; death means the user de-friends the brand, and the connection is lost.

The Risks of Quantification

This problem is greatly exacerbated by the current land rush to quantify the success of social media, which inevitably harms the cause of promoting quality. Consider another hypothetical scenario: a marketing manager convinces her VP that the brand could deepen their customer engagement using Twitter. The VP agrees, but demands that the manager establish a goal of 1,000 Twitter followers in the first quarter. Since the process of acquiring Twitter followers is largely organic, i.e., consumers encounter the brand's Twitter feed and decide it is

worth following, then the only way for the marketer to achieve this rapid-results goal is to offer an incentive for sign-ups through an online promotion.

The promotion succeeds in garnering the requisite Twitter follower count, but because these users signed up mainly to get the incentive, their engagement level is low. The marketer, in turn, is under pressure to monetize the channel, so she tries to counter low responsiveness with more frequent offers, which causes annoyed users to “un-follow” the brand. The VP concludes that social media has failed the brand, when, in fact, the brand has failed social media.

In this scenario, the Twitter strategy failed because it contained perverse incentives, i.e., it induced consumers to sign up for reasons that were at variance with the brand’s actual goal for being on Twitter. Consumers signed up for a prize or a discount irrespective of their long-term interest in connecting with the brand, and so the brand traded the quality of its followers for a quantity of followers that would look attractive on paper. The ultimate cost to the brand was far greater than if it had never been on Twitter at all.

Lest it appear that I am throwing the quantitative baby out with the perverse incentive bathwater, let me acknowledge that it is entirely reasonable for marketers to find ways to quantify the effects of social media, and it is self-evidently reasonable to set a goal of increasing one’s quantity of social media followers. In an economic downturn, the ability to quantify results is often what gets a marketing initiative funded, and funding social media initiatives can bring marketers closer to the 4-4 equilibrium.

The danger lies in treating quantification in absolute rather than relative terms, e.g., an arbitrary goal of x followers rather than gradually building a following over time. The absolutist approach is what leads to the problem of overgrazing; one can privilege quantity or quality, but never both at the same time. Hardin makes this point emphatically, citing game theory founder John von Neumann: “It is not mathematically possible to maximize for two (or more) variables at the same time.”

I will take the risk of belaboring this point because many marketers in the current climate seem determined to do what is mathematically impossible. The rush toward monetization of social media marketing creates perverse incentives for overgrazing; a case in point is the recent emergence of paid tweeting, i.e., the practice of paying an influential social networker to tout a product on Twitter in the context of their everyday tweeting. A *New York Times* article on the practice describes one influential Twitter user with 50,000 followers being paid to endorse personalized M&Ms candy (Stone 2009). This is a textbook example of trying to maximize both variables: the influencer is prized by the marketer for his/her large following, which is the

The absolutist approach is what leads to the problem of overgrazing; one can privilege quantity or quality, but never both at the same time.

result of their authenticity and credibility. But paid tweeting diminishes their authenticity and credibility, and once the payola is revealed or even suspected, the influencer will lose followers and credibility – both the quantity and the quality of the engagement are diminished.

But as with the overgrazers in Hardin’s metaphor, a single act of defection is easy to justify, because the single act alone will not ruin Twitter’s credibility, and the cost of its diminished credibility is born by the whole community, not by the solo defector. That defector cannot be blamed when the problem has become chronic across the network, and users abandon in droves.

It has been argued that as long as sponsors provide “transparency,” i.e., that they acknowledge when a message is sponsored, that users will not abandon the venue; they can simply choose to ignore or discount the message as they would any other advertising. But then what, actually, is the point?

If social connections between marketers and consumers bear the promise of a mutually beneficial equilibrium, but they are instead used like any other ad medium, what has either player gained? We would, in fact, see the same continuous performance decline in click-based Twitter or Facebook sponsorships as we did in banner advertising, in a cycle of mutual defection. In the *New York Times* piece, an owner of one of the social media sponsorship companies denied that sponsorships would diminish trust in social network connections, saying, “All we are trying to do is get consumers to become marketers for us.”

Therein lies the problem. Consumers are not marketers. As this book has tried to show, consumers and marketers are mutually dependent adversaries in the marketing game; each cannot do what the other does. When consumers naturally enthuse about brands in social media, their credibility comes from the fact that they are not marketers. Turning them into paid marketers destroys their credibility and ruins the system.

A marketer who read the *New York Times* article on Twitter with the slightest degree of hindsight would probably recognize the potential for ruination through overgrazing, because the problem is strongly reminiscent of the over-saturation of banner advertising in the name of quantification in the late 90’s, and the resultant consumer defection. But then what might the marketer conclude on the basis of this insight? I suggest there are two plausible responses:

1) “If this trend continues, Twitter sponsorships won’t be viable for very long. I had better take advantage now, while consumers are still responsive.”

Consumers are not marketers. As this study has tried to show, consumers and marketers are mutually dependent adversaries in the marketing game; each cannot do what the other does.

2) “If this trend continues, Twitter sponsorships won’t be viable for very long. I’m going to refrain from contributing to the problem.”

What I wish to point out is that neither of these responses solves the problem. The first response is obviously an outright defection that will accelerate the decline of the medium as surely as pop-under ads did in banner advertising. The second response is noble, cooperative, and forward-looking in its view that restraint is necessary to preserve the long-term good, but self-restraint will not prevent others from abusing the system. In order to prevent overgrazing of social media marketing, there must be behavioral controls in place that go beyond individual restraint.

Social Traps and Counterreinforcers in Social Media

The question of what kinds of behavioral controls would prevent overgrazing is the subject of “Social Traps,” a

In any online
community, implicit
rules of engagement
spring up very quickly,
and they tend to be
rigorously reinforced by
its membership.

breakthrough study by John Platt that analyzes the problem of reconciling individual self-interest to the collective good from the perspective of behavioral psychology. Platt builds on the groundwork laid by Garrett Hardin and Thomas Schelling in analyzing the volunteer’s dilemma, and he brings in Skinnerian behavioral psychology’s emphasis on how positive or negative behaviors are reinforced. In Platt’s simple and compelling formulation, social traps occur when a given behavior produces positive results for the individual and negative results for the group. As long as the individual is only accountable to themselves, the negative behavior is self-reinforcing, resulting in “locked-in behavior,” even though the individual’s long-term interests are imperiled by the behavior. This occurs because social traps typically involve a delay

between the short-term gain and the long-term loss; a farmer may get several years of good grazing from the commons before it is destroyed. This problem of “individual goods and collective bads” can’t be solved by the sacrifice of one or two heroes; positive group behavior must somehow be enforced (Platt 1973).

Platt offers several ways out of the social trap, some of which are applicable to the current social media marketing dilemma. The most important of these is the notion of “counterreinforcers.” Since destructive behavior is self-the player must evaluate before taking the action. For instance, if the herdsmen on the commons instituted a fee for every grazing animal added beyond a certain quota, then any herdsman acting in his short-term interest would have to weigh this cost against the profitability of adding another animal.

By social media’s very nature, the formation of such formal rules of collective engagement is rare, but counter-reinforcement is not. In any online community, implicit rules of engagement spring up very quickly, and they tend

to be rigorously reinforced by its membership. You'll recall the analysis of the psychological rewards of punishing bad behavior covered in Chapter 4; social media allows participants to go the extra mile in doling out punishment – particularly in the form of verbal castigation – at very little cost. Those who participated in the piling-on of negative reactions to the Motrin video paid very little: an investment of less than 5 minutes in viewing the offending video and responding on Twitter.

In this respect, the risks of a death spiral in social media are very different than they are in, say, banner advertising, where the marketer's level of control allows them to heap on more and more ad impressions, chasing the elusive click. In social media marketing, consumers exert great control over their level of exposure to brands and can easily dole out punishment in a variety of ways: negative feedback, de-friending, etc. It has often been noted that consumers in social media "vote with their feet," i.e., they quickly and easily drop social connections with brands that don't pay off for them.

Unlike a common grazing ground, social media is a common ground with a built-in set of constraints against bad behavior.

We could imagine, for instance, celebrities and/or influencers who engage in sponsored tweets being "un-followed" if the implicit rules of the community decreed that sponsored tweets are obnoxious and unwanted. Fearing for the integrity of their own personal brand in the face of this counter-reinforcement, the celebrity would have a strong incentive to drop the sponsorship. The consumer facing an unwanted marketing intrusion into, say, their movie-watching experience largely stands alone; their negative reaction never surfaces as a counter-reinforcer, and so the bad behavior continues. But an individual negative reaction on Twitter can be mustered into collective outrage in a matter of minutes. Thus unlike a common grazing ground, social media

is a common ground with a built-in set of constraints against bad behavior; it's very easy to get kicked of the collective. Marketers have a greater incentive to play by the rules.

Voting with their Feet: Why Quality Matters in Social Media

But marketers' implicit agreement not to be obnoxious is a rather low bar to set for a medium that offers opportunities for deep engagement, and indeed, winning the attention game will require more than avoiding bad behavior. Competing for a consumer's declining attention on a social network demands an emphasis on quality. If a consumer's increasingly divided attention span on Twitter means that they will only follow a handful of brands that provide them with valuable content, then there is a built-in incentive for brands to solve the problem with higher-quality Twitter content. If a consumer will only watch 1 out of every 100 brand-sponsored videos on YouTube, then quality, as measured in votes and popularity, will be the deciding factor. If the consumer's attention further subdivides to the point at which they will only watch 1 out of every 1000 brand videos, then quality must

increase accordingly. The loss of attention raises the cost of the signal that brands rely on to connect through social media, so that only brands able to pay the social cost of deep engagement will succeed. Brands that persist in the “something for nothing” view of social media will simply be squeezed out as available attention declines.

Given the social media acceleration I have outlined, it is axiomatic that consumers will become much more selective about their brand engagements in social media, even as most brands are still getting their sea legs. Does this mean that some brands will simply fail at social media? *It does*. In fact, given the potential dangers of saturation, it is *necessary* that some brands fail at social media, so that others can succeed. Defining success on the basis of quality rather than quantity is social media marketing’s best chance at becoming a mature marketing medium.

Because analysis of social media marketing is still in its Unbridled Enthusiasm stage, in which every new venue is treated as the next big thing, very little attention has been paid to the qualitative factors that determine a brand’s success in the space. As I have already noted, the popularity-based system of costly signaling will allow some brands to succeed at lower cost than others. But how does social media separate the wheat from the chaff? The best work recognizing the importance of attention as a limited commodity in the social media game is an overlooked study from HP Laboratories’ Social Computing Lab, “Crowdsourcing, Attention and Productivity.” The study’s subject is the relationship between popularity and productivity in YouTube videos, but its findings are broadly applicable to social media marketing. The study’s authors raise the question of whether a “tragedy of the commons” is unfolding on YouTube, where over-competition for user attention discourages users from producing new content. In marketing terms, this would impact both brands’ willingness to provide content and users’ willingness to produce their own brand-related content (such as the Starbucks fan’s chronicle of his efforts to visit every Manhattan Starbucks).

The study found that attention was indeed the valued commodity that YouTube uploaders pursued, independent of financial gain. The attention paid by other users, measured in views and comments, very strongly correlated to the likelihood that users would produce more content, and lack of attention had the inverse effect, to the point where users that lacked attention would stop producing videos (Huberman 2009). In other words, the factor that hedges against oversaturation and the tragedy of the commons is the phenomenon of users “voting with their feet”; contributors and brands that don’t achieve good quality scores – in the forms of views, comments, and votes – will decide that the costly signal

The loss of attention raises the cost of the signal that brands rely on to connect through social media, so that only brands able to pay the social cost of deep engagement will succeed.

of popularity is too dear. The resultant equilibrium will indeed exclude some brands and force all participating brands to work harder to gain popularity, but the tragedy of the commons can be averted.

Dear Reader,

You made it this far – is it possible you'd like more? In other chapters, I put game theory to work on such tantalizing marketing topics as:

- How zero-sum theory can improve banner advertising
- How nuclear détente predicted social media marketing
- Why consumers need to punish marketers on YouTube
- How Coca-Cola bagged a stag on Facebook
- How Skittles got burned by the Prisoner's Dilemma
- How Starbucks took a page from the Roman army's playbook
- Why most viral videos need to fail

All of this and much more is coming your way in *Social Media Marketing: A Game Theory Perspective*.

For updates on the book's release, follow me on Twitter @unsettler, or email me at eanderson@whitehorse.com . Look for the book on Amazon and Springer.com in June 2010.

Eric Anderson | White Horse | eanderson@whitehorse.com